AOARD Grant: 104123
Learning within Optimization
Final Report

Peter Stuckey, Toby Walsh

NICTA, Australia
email: {peter.stuckey,toby.walsh}@nicta.com.au

Final Report

This report summarizes research undertaken for AOARD Grant 104123 awarded to Prof. Peter Stuckey (Melbourne, Victoria) and Prof. Toby Walsh (Sydney, New South Wales).

Summary

Combinatorial optimization problems are ubiquitous and important. Such problems arise anywhere a set of interlinking decisions need to be made. For example, in a transportation problem, we might want to determine which trucks will deliver which goods in which order to minimize costs. Learning is an important but neglected aspect of combinatorial optimization. We want our solver to learn how to solve such optimization problems. In addition, as the world is dynamic and uncertain, we want our solver to learn how the constraints of our problem change. The goal then of this project was to develop a theoretical and practical understanding of how we can combine learning within optimization. In particular, we developed a powerful learning based search method called lazy-clause generation. This has now been shown to advance the state of the art in several domains, especially with the optimisation subfield of scheduling.

Personnel

In New South Wales, Dr Nina Narodytska was hired at the start of the grant as a postdoc, funded in part by AOARD and the rest by NICTA. In Victoria, Dr Andreas Schutt was appointed as a postdoc also funded in part by AOARD.

Additional information

There were no significant changes in research objectives. No extensions were granted or milestones slipped.

Report Documentation Page

Form Approved
OMB No. 0704-0188

1. REPORT DATE	2. REPORT TYPE	3. DATES COVERED
18 APR 2013	**Final**	**01-09-2010 to 31-08-2012**

4. TITLE AND SUBTITLE	5a. CONTRACT NUMBER
Learning Within Optimization	**FA23861014123**
	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)	5d. PROJECT NUMBER
Toby Walsh	5e. TASK NUMBER
	5f. WORK UNIT NUMBER
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)	8. PERFORMING ORGANIZATION REPORT NUMBER
National ICT Australia United,Locked Bag 6016,Sydney, NSW 2052,Australia,NA,NA	**N/A**
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR'S ACRONYM(S)
AOARD, UNIT 45002, APO, AP, 96338-5002	**AOARD**
	11. SPONSOR/MONITOR'S REPORT NUMBER(S)
	AOARD-104123

12. DISTRIBUTION/AVAILABILITY STATEMENT

Approved for public release; distribution unlimited

13. SUPPLEMENTARY NOTES

14. ABSTRACT

Combinatorial optimization problems are ubiquitous and important. Such problems arise anywhere a set of interlinking decisions need to be made. For example, in a transportation problem, we might want to determine which trucks will deliver which goods in which order to minimize costs. Learning is an important but neglected aspect of combinatorial optimization. We want our solver to learn how to solve such optimization problems. In addition, as the world is dynamic and uncertain, we want our solver to learn how the constraints of our problem change. The goal then of this project was to develop a theoretical and practical understanding of how we can combine learning within optimization. In particular, we developed a powerful learning based search method called lazy-clause generation. This has now been shown to advance the state of the art in several domains, especially with the optimisation subfield of scheduling.

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			
unclassified	**unclassified**	**unclassified**	**Same as Report (SAR)**	**3**	

Detailed research description

Lazy clause generation

One way to make optimisation more effective is to learn *during* search. Lazy clause generation is an approach to learning during search by recording the reasons that some previous part of the search failed in order to avoid repeating the same failed search later on. This prize winning method is highly effective in solving many combinatorial optimization problems. We continue to investigate improvements to lazy clause generation, as well as it application to other optimisation techniques like answer set programming.

C. Drescher, T. Walsh: Answer Set Solving with Lazy Nogood Generation. Technical Communications of the 28th International Conference on Logic Programming, ICLP 2012: 188-200.

A. Schutt, G. Chu, P.J. Stuckey, and M. Wallace. Maximizing the net-present-value for resource constrained project scheduling. In N. Beldiceanu, N. Jussien, and E. Pinson, editors, International Conference on Integration of Artificial Intelligence and Operations Research Techniques in Constraint Programming for Combinatorial Optimization Problems (CPAIOR), number 7298 in LNCS, pages 362-378. Springer, 2012.

This last paper defines the state of the art complete method for solving resource constrained scheduling problems where the aim is to maximize net-present value. In such scheduling problems tasks are associated with a cost or benefit which evolves with a fixed interest rate, so costly tasks should be delayed as long as possible, and beneficial tasks completed as soon as possible. By combining learning approaches we were able to solve the benchmark suite for this challening optimization problem an order of magnitude faster than the previous state of the art.

Another challenge is to learn across problem instances. The optimisation problem we solve today is often closely related to the one we solved yesterday, and the one we will be asked to solve tomorrow.

G. Chu and P.J. Stuckey. Inter-problem nogood learning in constraint programming. In M. Milano, editor, Proceedings of the 18th International Conference on Principles and Practice of Constraint Programming, number 7514 in LNCS, pages 238-247. Springer, 2012.

This paper shows how we can reuse what we have learned in solving one optimization problem for solving a similar optimization problem. In that, it is an important first step in the research of lifelong optimization, since we expect an repeated optimization problem to change slowly over its lifetime, being able to take advantage of earlier learnt information (in this paper learnt clauses) can give significantly improve solving of later iterations of the lifelong optimization problem.

Symmetry breaking

Another significant research challenge is to learn and automate methods for modelling problems. One area that we have focused on here is dealing with symmetries.

T. Walsh: Symmetry Breaking Constraints: Recent Results. Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence. AAAI Press 2012.

One type of symmetry in optimsation problems is dominance. We often have a situation where one solution dominates another. That is, one essentially symmetric solution is better (cheaper or faster) than another.

G. Chu and P.J. Stuckey. A generic method for systematically identifying and exploiting dominance relations. In M. Milano, editor, Proceedings of the 18th International Conference on Principles and Practice of Constraint Programming, number 7514 in LNCS, pages 6-22. Springer, 2012.

This paper shows gives a generic method for creating dominance breaking constraints for arbitrary optimizaiton problems. Adding dominance breaking constraints has up until this paper been an ad hoc affair, with new methods for each problem. This generic approach is compatible with any form of search, and thus in particular most valuable when we are combining learning with optimization.

Awards

Peter Stuckey's work on dominance was recognized with the Best Paper award at the 18th International Conference on Principles and Practice of Constraint Programming in October 2012.

Peter Stuckey's work on lazy clause generation won the inaugural Google Australia Eureka Prize for Innovation in Computer Science. The prize is the highest award in computer science in Australia. He was also awarded the University of Melbourne's Woodward Medal for Science and Technology for this work.

Toby Walsh and Nina Nina Narodytska received the Outstanding Paper Award at the Twenty-Fifth Conference on Artificial Intelligence (AAAI-11) in San Francisco, California during August 2011.

Toby Walsh received the Best Poster Award at Thirteenth Conference on Theoretical Aspects of Rationality and Knowledge (TARK XIII) in Groningen, the Netherlands during July 2011.